Close-range Photometric Stereo with Point Light Sources

Aaron Wetzler, Ron Kimmel, Alfred M. Bruckstein
Computer Science Department,
Technion, Israel Institute of Technology,
Technion City, Haifa 32000, Israel
twerd,ron,freddy@cs.technion.ac.il

Roberto Mecca
Pattern Analysis & Computer Vision Department,
Istituto Italiano di Tecnologia,
via Morego, 30, 16163 Genova
roberto.mecca@iit.it

***Abstract*—Shape recovery based on shading variations of a lighted object was recently revisited with improvements that allow for the photometric stereo approach to serve as a competitive alternative for other shape reconstruction methods. However, most efforts of using photometric stereo tend to ignore some factors that are relevant in practical applications. The approach we consider tackles the photometric stereo reconstruction in the case of near-field imaging which means that both camera and light sources are close to the imaged object. The known challenges that characterize the problem involve perspective viewing geometry, attenuation of light and possibly missing regions. Here, we pay special attention to the question of how to faithfully model these aspects and by the same token design an efficient and robust numerical solver. We present a well-posed mathematical representation that integrates the above assumptions into a single coherent model. The surface reconstruction in our near-field scenario can then be executed efficiently in linear time. The merging strategy of the irradiance equations provided for each light source allows us to consider a characteristic expansion model which enables the direct computation of the surface. We evaluate several types of light attenuation models with nonuniform albedo and noise on synthetic data using four virtual sources. We also demonstrate the proposed method on surface reconstruction of real data using three images, each one taken with a different light source.**

***Keywords*-Photometric Stereo; Partial Differential Equations; Shape from Shading**

I. Perspective On Shape from Shading and Photometric Stereo

Since the seminal Shape from Shading paper by B.K.P. Horn [11], new models have been introduced in order to extend the range of shape recovery problems that can be solved by Shape from Shading (SfS) methods [16], [21], [22]. A particularly important direction of research has been the transition from the assumption of orthographic viewing geometry [13], [4], [12] to the more realistic one of perspective cameras [19].

Several papers reported using Perspective SfS (PSfS) methods applied to endoscopic image analysis [22], [21], [16]. Deguchi et al. [16] introduced the perspective viewing shape estimation for objects located close to the camera, taking into account a realistic endoscopic model for close light source illumination (i.e. ideally attached to the optical center). Later, Prados et al. [19] studied the same formulation under a *differential* point of view and, using a model based on Partial Differential Equations (PDEs), they concluded that their formulation for the endoscopic perspective shape from shading problem is well-posed. A recent paper by Breuß et al. [3] shows that there might be ambiguity in the Prados et al. model. In fact, near-field Endoscopic Perspective (EP) shape reconstruction remains an open problem if only a single image is considered.

In this paper we deal with the case of calibrated Photometric Stereo (PS), where multiple images of the same scene are taken under different illumination conditions (with known light positions) when scanner and object are close to each other. For a recent overview about calibrated (and uncalibrated) PS methods see [17]. In the EP framework Wu et al. [21] studied the multi-image EP problem by considering two light sources placed off the optical center. They develop a model based on the radiance information obtained by simultaneously illuminating an object with two different light sources. They then recover the surface by considering a single irradiance equation for the sum of Lambertian reflectance functions of the two different light sources. The use of this reflectance function results in a loss of information. In order to avoid this problem and issues related to unknown albedo, they use a photometric calibration. Surface recovery is performed within a variational framework that involves high computational complexity compared to alternative direct methods [14]. The shape from EP problem solved via a PS technique using n images (EPPS$_n$) is a problem initially addressed by Iwahori et al. [**?**], Clark [**?**] and more recently Migita et al. [**?**]. Collins and Bartoli [7] solve the close-range PS problem with weak considerations on the parametrization of the variables like unknown surface, light sources and solving the problem of the light attenuation with a-priori light calibration procedure. They also use a prior for a reflectance model, adding patches on the inspected object even if the surface is assumed to be Lambertian. In particular, their mathematical formulation is based on the usual double step procedure where an energy functional is minimized (which allows the computation of the surface derivatives) and thereafter the surface is recovered [9], [1]. Let us emphasize that their energy is based on the sum of Lambertian irradiance equations and this is a sort of

contradictory procedure with respect to the scientific trend using photometric ratios [20], [14], [5] which yield more suitable problems. For example the ratio approach effectively makes the problem independent from the albedo.

Recently Parot et al. [18] studied the EPPS$_n$ by using the classic approach to PS and using a heuristic filtering process. They solve for the normal field to the surface as if the light sources were distant and then filter the directional gradients depending on the frequencies. The depth map is then computed using a multi grid solver for the Poisson equation. The applicability of this method is very limited since they calibrate the uniform light directions assuming *reasonable* distance between the object and the scanner. This allows to exaggerate surface features but not reconstruct the actual surface depths.

Here, we present a new mathematical formulation for the EPPS$_n$ by considering the endoscopic perspective presented in [16] and using the perspective parametrization used in [15] based on non-linear PDEs from photometric ratios. Such an extension introduces several realistic features which make the new model mathematically interesting. The formulation of a new mathematical model based on quasi-linear PDEs is followed by an efficient finite difference upwind numerical scheme for the direct approximation of the surface. This scheme is simple to implement, highly parallelizable and converges efficiently. The new model we describe can also successfully handle images with missing data. Usually the image acquisition procedure for PS includes nuisance factors such as occlusions and shadows. Here we do not focus on occlusions or shadows detection, but we show that our problem successfully reconstructs surfaces when images have missing parts. Furthermore, a significant advantage of our model is the direct computation of the 3D surface without explicit computation of the surface normals for the purpose of integration. Non-linearities resulting from our realistic lighting model yield a normal field that depends on the depth of the surface as well as the direction of the light source. However, this does not pose a problem in our framework because we avoid directly computing the normal field and using it in the numerical scheme.

In Section II we recall a well known parametrization for the surface under close camera observation. The new mathematical model is introduced in Section III where we show the differential model including non linear light attenuation. Section IV introduces the theoretical formulation of the new differential approach for three images which can be easily extended to $n \geq 3$. Using this formulation we will show how it is possible to overcome the problem of handling images with missing regions. In Section V we solve the problem of EPPS$_3$ and EPPS$_4$ when there are images with missing parts and an unknown albedo. A portion of that section is also devoted to the explanation of the up-wind scheme used in the numerical tests. Finally, in Section VI the experimental results on synthetic and real data are presented.

II. The Endoscopic Perspective Set-Up

In order to give all the necessary ingredients to understand the geometry behind the model we start by considering the parametrization of the surface Σ (see Fig. 1) given in [19] up to an unknown function z from the image domain $\overline{\Omega}_p = \Omega_p \cup \partial\Omega_p$ to $\mathbb{R}$ such that:

$$\mathbf{M}(x,y) = [\xi(x,y), \eta(x,y), \zeta(x,y)] := \left[-x\frac{z(x,y)}{f}, -y\frac{z(x,y)}{f}, z(x,y) \right]. \quad (1)$$

Here, $f > 0$ is the focal length of the camera, $\zeta < -f < 0$, and the triple $[\xi(x,y), \eta(x,y), \zeta(x,y)] = [\xi, \eta, \zeta]$ are the real world coordinates (with respect to the image coordinates). This parametrization (1) is based on the pinhole camera model and is due to the specific perspective viewing geometry as seen in Fig. 1 where the camera is placed at the origin C of the coordinate system $C\xi\eta\zeta$ (namely at the optical center) [19].

We recall that an outgoing vector normal to the surface Σ is

$$\overline{\mathbf{n}}(x,y) = \frac{z}{f^2}\big[f\nabla z(x,y), z(x,y) + (x,y)\cdot\nabla z(x,y)\big], \quad (2)$$

and since the irradiance equation depends on the unit normal, we take $\mathbf{n}(x,y) = \frac{\overline{\mathbf{n}}(x,y)}{|\overline{\mathbf{n}}(x,y)|}$ as a unit length normal vector.

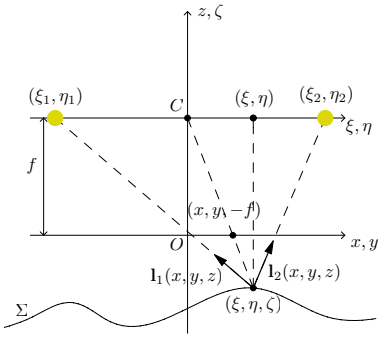

Figure 1. In the perspective world (i.e. the image coordinate system given by $Oxyz$) the light reflected at image point (x,y) comes form the real point $[\xi,\eta,\zeta]$ of the surface.

We consider the well-known irradiance equation for Lambertian surfaces, given by the cosine law by the following inner product

$$I(x,y) = \rho(x,y)(\mathbf{n}(x,y)\cdot\mathbf{l}(x,y,\ldots)), \quad (3)$$

where $I : \overline{\Omega}_p \to [0,1]$ is the image function, $\rho(x,y)$ is the unknown albedo and $\mathbf{l}(x,y,\ldots)$ is the light source direction. Let us emphasize that the light direction $\mathbf{l}$ incident to the surface depends explicitly on the image points (x,y) and on other factors $(\ldots)$ we shall specify in the next section, since the endoscopic formulation assumes close light source illumination.

III. A New Close-Range PS Model

In much of the literature dealing with Perspective SfS [19], [16], a single light source is considered to be placed at the optical center. This assumption is unrealistic when considering a camera close to the inspected object. Since PS uses several light sources, we consider a considerably more realistic placement (ξ_j, η_j) on the optical plane, that is $\zeta = 0$. The coplanarity of the sources is not a necessary constraint but we adopt it here to simplify our formulation. We define the light directions as

$$\mathbf{l}_j(x,y,z) = \frac{\left[-\frac{\xi_j f}{z} - x, -\frac{\eta_j f}{z} - y, f\right]}{\sqrt{\left(x + \frac{\xi_j f}{z}\right)^2 + \left(y + \frac{\eta_j f}{z}\right)^2 + f^2}} = \frac{\bar{\mathbf{l}}_j(x,y,z)}{q_j(x,y,z)}, \tag{4}$$

where

$$q_j(x,y,z) = \sqrt{\left(x + \frac{\xi_j f}{z}\right)^2 + \left(y + \frac{\eta_j f}{z}\right)^2 + f^2}. \tag{5}$$

Let us emphasize that for this model, the light directions depend not only on the point (x, y), but since they are displaced from the optical center, they also depend on z. This introduces a non-linearity that does not involve the derivatives of z.

A. Two types of light attenuation

We now consider two different kinds of light attenuation. The first factor is due to the reduction of light energy proportional to the inverse squared distance between the light source and object. The second factor of attenuation we describe, is a result of a realistic directional lighting model of a real SMD light. However in principle we could choose among many different continuous attenuation models to suit our needs.

1) Light attenuation due to distance: The standard way to model this attenuation of the light intensity is to compute the distance between the light source and the surface explicitly. For this purpose, let us consider the functions $\mathbf{r}_1(x,y,z)$ and $\mathbf{r}_2(x,y,z)$ as the distances between the point of the surface $[\xi, \eta, \zeta]$ and the respective light source.

In this case, since the light sources are shifted with respect to the origin, the distance function is as follows:

$$r_j(x,y,z) = \text{dist}\left((\xi_j, \eta_j, 0), \left(-x\frac{z}{f}, -y\frac{z}{f}, z\right)\right) = \frac{z}{f}\sqrt{\left(x + \frac{\xi_j f}{z}\right)^2 + \left(y + \frac{\eta_j f}{z}\right)^2 + f^2} = \frac{z}{f} q_j(x,y,z). \tag{6}$$

The attenuation factor can then be written as r_j^{-2}.

2) Radial attenuation of the light: Many existing light sources are directional. That is to say, they are bright along a principal direction and become less bright at angles further from the main direction. This behavior can be observed in Fig. 2 and can be effectively simulated by multiplication with $\cos(\theta)^\mu$ where μ is the attenuation coefficient and is reminiscent of the specular model for surface reflectance.

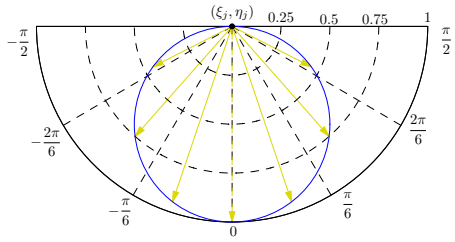

Figure 2. The intensity of light for a directional light source pointing downwards.

The attenuation factor is easily computable since

$$\cos(\theta)^\mu(x,y,z) = (\mathbf{l}_j(x,y,z) \cdot (0,0,1))^\mu = \frac{f^\mu}{q_j^\mu(x,y,z)} \tag{7}$$

where $\mathbf{l}_j(x,y,z)$ is the i-th light source placed at (ξ_j, η_j). The light direction is collinear with the viewing direction of the camera which simplifies both our formulation and the experimental setup. Both attenuation effects can be expressed by multiplication of the following factor

$$\mathrm{a}_j(x,y,z) = \frac{f^\mu}{r_j^2(x,y,z) q_j^\mu(x,y,z)} = \frac{f^{\mu+2}}{z^2 q_j^{\mu+2}(x,y,z)}. \tag{8}$$

B. EPPS$_2$ model

Our model for the SfEPPS$_2$ problem considers the following irradiance equations

$$\begin{aligned} I_j(x,y) =& \rho(x,y)\frac{\bar{\mathbf{l}}_j(x,y,z)\cdot\mathbf{n}(x,y)}{r_j^2(x,y,z) q_j^{\mu+1}(x,y,z)} f^\mu = \\ & \rho(x,y)\frac{\bar{\mathbf{l}}_j(x,y,z)\cdot\overline{\mathbf{n}}(x,y) f^{\mu+2}}{z^2 q_j^{\mu+3}(x,y,z)|\overline{\mathbf{n}}(x,y)|} \end{aligned} \tag{9}$$

where

$$\bar{\mathbf{l}}_j(x,y,z)\cdot\overline{\mathbf{n}}(x,y) = -\xi_j z_x - \eta_j z_y + \frac{z^2}{f}. \tag{10}$$

Now, solving the EPPS$_2$ from a mathematical point of view consists of solving the following system of non-linear PDEs of Hamilton-Jacobi type,

$$\begin{cases} I_1(x,y) = \rho(x,y)\dfrac{\bar{\mathbf{l}}_1(x,y,z)\cdot\overline{\mathbf{n}}(x,y) f^{\mu+2}}{z^2 q_1^{\mu+3}(x,y,z)|\overline{\mathbf{n}}(x,y)|} & \text{on } \Omega_p \\ I_2(x,y) = \rho(x,y)\dfrac{\bar{\mathbf{l}}_2(x,y,z)\cdot\overline{n}(x,y) f^{\mu+2}}{z^2 q_2^{\mu+3}(x,y,z)|\overline{\mathbf{n}}(x,y)|} & \text{on } \Omega_p \\ z(x,y) = g(x,y) & \text{on } \partial\Omega_p \end{cases}. \tag{11}$$

Our strategy to solve such a problem is to merge the irradiance equations of (11) by noting that the non-vanishing quantity $\frac{\rho(x,y)f^{\mu+2}}{|\overline{\mathbf{n}}(x,y)|z^2}$ is present in both of them. We merge such equations as follows

$$\underbrace{\frac{I_1(x,y)q_1^{\mu+3}(x,y,z)}{\bar{\mathbf{l}}_1(x,y,z)\cdot\overline{\mathbf{n}}(x,y)}}_{\text{First Equation of (11)}} = \overbrace{\frac{\rho(x,y)f^{\mu+2}}{|\overline{\mathbf{n}}(x,y)|z^2}}^{\text{Second Equation of (11)}} = \frac{I_2(x,y)q_2^{\mu+3}(x,y,z)}{\bar{\mathbf{l}}_2(x,y,z)\cdot\overline{\mathbf{n}}(x,y)}$$

getting

$$\frac{I_1(x,y)q_1^{\mu+3}(x,y,z)}{-\xi_1 z_x - \eta_1 z_y + \frac{z^2}{f}} = \frac{I_2(x,y)q_2^{\mu+3}(x,y,z)}{-\xi_2 z_x - \eta_2 z_y + \frac{z^2}{f}}. \quad (12)$$

Let us emphasize that this way of merging the irradiance equations, allows us to eliminate the non-linearity with respect to the partial derivatives of z contained in $|\overline{\mathbf{n}}(x,y)|$. This makes the resulting problem completely independent from the albedo and the strategy is an incremental development over [20], where the ratio idea originated.

After some algebra, we obtain the following first-order quasi-linear PDE

$$\begin{aligned} &z_x\Big(I_1(x,y)q_1^{\mu+3}(x,y,z)\xi_2 - I_2(x,y)q_2^{\mu+3}(x,y,z)\xi_1\Big) + \\ &z_y\Big(I_1(x,y)q_1^{\mu+3}(x,y,z)\eta_2 - I_2(x,y)q_2^{\mu+3}(x,y,z)\eta_1\Big) = \\ &\frac{z^2}{f}\Big(I_1(x,y)q_1^{\mu+3}(x,y,z) - I_2(x,y)q_2^{\mu+3}(x,y,z)\Big), \quad (13) \end{aligned}$$

resumed as follows

$$\begin{cases} \mathbf{b}(x,y,z)\cdot\nabla z(x,y) = s(x,y,z), & \text{on } \Omega_p \\ z(x,y) = g(x,y) & \text{on } \partial\Omega_p, \end{cases} \quad (14)$$

where $g(x,y)$ is the Dirichlet boundary condition,

$$\begin{aligned} \mathbf{b}(x,y,z) = \Big[&I_1(x,y)q_1^{\mu+3}(x,y,z)\xi_2 - I_2(x,y)q_2^{\mu+3}(x,y,z)\xi_1, \\ &I_1(x,y)q_1^{\mu+3}(x,y,z)\eta_2 - I_2(x,y)q_2^{\mu+3}(x,y,z)\eta_1\Big] \end{aligned}$$

and

$$s(x,y,z) = \frac{z^2}{f}\Big(I_1(x,y)q_1^{\mu+3}(x,y,z) - I_2(x,y)q_2^{\mu+3}(x,y,z)\Big).$$

IV. Direct Surface Reconstruction Using Images with Missing Parts

If we have three images then we can consider the set of unique image pairs and linearly combine them following the same procedure as [15]. We are thus able to define the EPPS$_n$ problem by exploiting the linearity of the basic differential formulation (14) and reducing it to a single quasi-linear PDE which can handle missing regions in a natural fashion.

Since we want to exploit the PS technique, we assume that each pixel is illuminated in at least two images thereby avoiding reduction to a EPPS$_1$ problem. This means that we need a minimum of two light sources to proceed with our method.

In order to complete the theoretical analysis, we note that the uniqueness of the weak (i.e. Lipschitz, $z \in \text{Lip}(\overline{\Omega}_p)$) solution for the differential problem (14) can be proved ([14]). Here an important comment must be made regarding the light attenuation. In the previous section we introduced the new model (14) with two specific light attenuations because for real tests we used SMD lights with these photometric characteristics. However, well-posedness is guaranteed even for general attenuation functions (8). In the following sections we show how we can still reconstruct the surface from three images, even if they have non-overlapped missing parts, treating the new differential problem with a general attenuation function $a_j(x,y,z)$.

V. W-EPPS$_n$ with Almost no Boundary Condition

In this section we focus on the applicability of our model. In Section IV we extended the EPPS$_n$ model by supposing the knowledge of the boundary condition $g(x,y)$. Clearly such a hypothesis compromises the use of that model for many real applications. It is therefore important to find a way to solve the EPPS$_n$ problem while removing the requirement for a-priori knowledge of the boundary condition.

Researchers facing the SfS, PSfS or the EPPS problems often adopt a two-stage strategy where one first computes the surface normals all over the domain and then integrates the normal field in order to perform surface recovery. Usually the normal computation is done by considering a linear system computed using the irradiance equations. Let us emphasise that when close light sources are taken into account the usual linearity of the Lambertian reflectance equation is lost. In this particular case, since the light sources are not placed at the optical center, the normalization of the light direction and the further light attenuations introduce a non-linearity with respect to z. In other words, given n images, it is very hard to compute the unknown z_x, z_y and z just by solving the following system of irradiance equations:

$$\begin{cases} I_1 = \rho(x,y)\mathrm{a}_1(x,y,z)\dfrac{\bar{\mathbf{l}}_1(x,y,z)\cdot\mathbf{n}(x,y)}{q_1(x,y,z)} \\ I_2 = \rho(x,y)\mathrm{a}_2(x,y,z)\dfrac{\bar{\mathbf{l}}_2(x,y,z)\cdot\mathbf{n}(x,y)}{q_2(x,y,z)} \\ \vdots \\ I_n = \rho(x,y)\mathrm{a}_n(x,y,z)\dfrac{\bar{\mathbf{l}}_n(x,y,z)\cdot\mathbf{n}(x,y)}{q_n(x,y,z)} \end{cases}. \quad (15)$$

The approach in this work is to compute z without completely eliminating the non-linearity. In fact, the quasi-linear PDE of (14) still contains a non-linear component, but the problem of recovering z can be successfully achieved directly by solving a quasi-linear PDE.

Furthermore, partial images clearly represent a loss of information in the image set and several authors have

approached the problem of surface recovery with occlusions [10], [6]. However we are not aware of attempts to overcome the problem of missing parts within the framework of the endoscopic problem.

We take the above issues into consideration and use them to help design a numerical strategy for reconstructing the surface using our model. The strategy involves determining the depth of only a single arbitrarily placed initial seed point within the reconstruction domain and robustly manipulating the path of the characteristics from that point. We do this in order to let the information travel in the most convenient directions for the whole domain.

A. Steering the characteristic field

On the way to defining a numerical strategy we will need to manipulate the path along which the information travels. To do this we will exploit the following result:

Theorem *Let $\mathbf{b}_t(x,y,z)$ be the vector field of any pair of images where $t \in \binom{[n]}{2}$[1]. Then, $\forall t_1, t_2 \in \binom{[n]}{2}$, $\forall (u,v) \in \Omega_p$ and $\forall z \in Lip(\overline{\Omega}_p)$ we have:*

$$\mathbf{b}_{t_1}(x,y,z) \cdot \mathbf{b}_{t_2}(x,y,z) \neq \pm |\mathbf{b}_{t_1}(x,y,z)||\mathbf{b}_{t_2}(x,y,z)|. \tag{16}$$

In other words, this theorem states that two different vector fields $\mathbf{b}_{t_1}$ and $\mathbf{b}_{t_2}$ can not be parallel. We can adopt the same fast marching strategy of [15] steering the direction of the characteristics for the case when $n = 3$.

We enumerate the steps in our characteristic steering method here:

1) fix the exact depth value to z for a point, in our case it shall be towards the center of the image domain adding all of that point's neighbors to a list of pixels to be visited;
2) traverse the list of pixels to be visited and update the value for z for each one by the scheme (18) derived in the next section;
3) for each newly visited pixel add its unvisited neighbors to the list of pixels to be visited;
4) in case of (non-overlapped) missing parts in the images, we can change the wavefront propagation direction in order to *surround* the shadow sets (i.e. computing the boundary condition) and then solve (14) with the appropriate pair of images;
5) the above steps are repeated until the L^∞ discrete norm of the difference between the last two elements of the approximating sequence is smaller than some predetermined threshold.

We remark that in this near-field set up, the convergence of the previous algorithm is not as straightforward as in [15] because here the vector fields $\mathbf{b}(x,y,z)$ depend on z, i.e. they are unknown. However, this aspect does not prevent the scheme from converging efficiently. The fast convergence behaviour of the problem requires a non-trivial proof showing the convergence of the numerical scheme with respect to [14], but we omit it here due to lack of space.

[1] $\binom{[n]}{2}$ is the set of pairs of integer indices with no repetition. For example, if $n = 3$, we have $\binom{[3]}{2} = \{(1,2),(1,3),(2,3)\}$.

B. Numerical scheme

We now consider the numerical method that we employ to obtain the experimental results. The scheme considered originates from [14] where a finite difference up-wind scheme is used.

In order to simplify the notation we shall denote $\mathbf{b}(x_i, y_j, z(x_i, y_j))$ by $\mathbf{b}_{i,j}(z_{i,j}) = (b^1_{i,j}(z_{i,j}), b^2_{i,j}(z_{i,j}))$ and $s(x_i, y_j, z(x_i, y_j))$ by $s_{i,j}(z_{i,j})$.

Let us consider the following implicit up-wind scheme:

$$\begin{aligned} b^1_{i,j}(Z_{i,j})\frac{Z_{i+1,j} - Z_{i-1,j}}{2\Delta} + b^2_{i,j}(Z_{i,j})\frac{Z_{i,j+1} - Z_{i,j-1}}{2\Delta} = \\ s_{i,j}(Z_{i,j}) + |b^1_{i,j}(Z_{i,j})|\frac{Z_{i+1,j} - 2Z_{i,j} + Z_{i-1,j}}{2\Delta} + \\ |b^2_{i,j}(Z_{i,j})|\frac{Z_{i,j+1} - 2Z_{i,j} + Z_{i,j-1}}{2\Delta}. \end{aligned} \tag{17}$$

The artificial diffusion introduced in the right side of (17) allows us to follow the vector field $\mathbf{b}$ by considering the most appropriate discretization for the fist derivative in order to track the characteristic lines. Specifically it consists of a numerical scheme of consistency order equal to one with respect to both partial derivatives.

We can write (17) explicitly as follows:

$$\begin{aligned} Z^{r+1}_{i,j} = \Big(|b^1_{i,j}(Z^r_{i,j})|Z^r_{i-\mathrm{sgn}\,(b^1_{i,j}(Z^r_{i,j})),j} + \\ |b^2_{i,j}(Z^r_{i,j})|Z^r_{i,j-\mathrm{sgn}\,(b^2_{i,j}(Z^r_{i,j}))} + s_{i,j}(Z^r_{i,j})\Delta\Big)\frac{1}{|b^1_{i,j}(Z^r_{i,j})| + |b^2_{i,j}(Z^r_{i,j})|}. \end{aligned} \tag{18}$$

VI. Experimental Results

We now describe the full experimental procedure we followed while investigating the properties of the new model. The numerical schemes were all implemented in Matlab MEX files in unoptimized C++ using OpenMP for the parallelization and executed in Matlab using a 2013 Dell Precision M6700 with an Intel i7 CPU clocked at 3GHz and 32GB of RAM.

A. Synthetic cases

1) Eve: We therefore first consider a realistic shape obtained from a mannequin head where real depth data has been obtained from a structured light range scanner. This is stored as a height field over a grid with 4-neighbor connectivity. The mannequin face (known as Eve) is approximately 70cm from the virtual camera center. The virtual camera is a perspective pinhole camera and all depth pixels which are background are set to NaN to indicate that they are not part of the mesh. Each pixel is also assigned an albedo which we restrict to be between 0 and 1. For the experiments with Eve we use either a uniform albedo of 1 or a synthetically

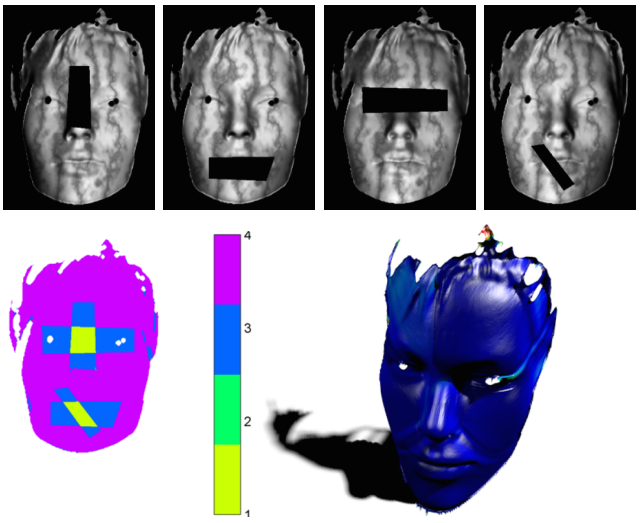

Figure 3. The input images with $\mu = 1$ are shown in the first row. The second row has (on the left) the shadow map and (on the right) the approximate 3D shape with the Euclidean error map textured on. $MSE = 0.52mm^2$

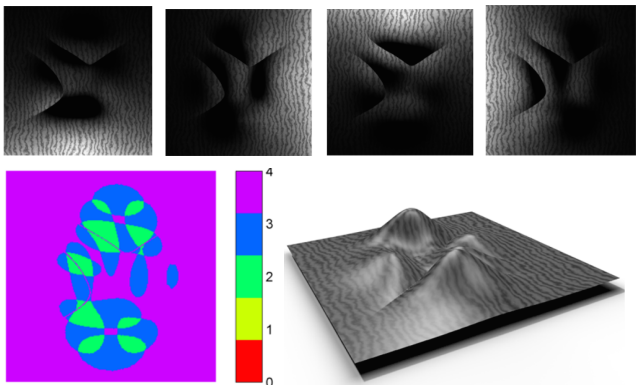

Figure 4. The input images with $\mu = 5$ are shown in the first row. The second row has (on the left) the shadow map and (on the right) the approximate 3D shape with recovered albedo textured on. $MSE = 3.75\text{E-}4units^2$

generated albedo based on Perlin noise. We define four virtual light sources which lie on the camera plane $z = 0$ and are positioned at 90^o intervals at a radius of 4cm. Each light source is defined by its direction and the non-linear light attenuation coefficient μ as described in Section III-A2.

Using the above setup we performed the experiment as shown in Fig. 3. Here the central pixel was initialized with the ground truth depth value and the method described in Section V-A was used together with the backward upwind scheme of Section V-B.

2) AbsPeaks, with realistic shadows: Fig. 4 shows the set of significantly overlapped shadows and light attenuations where $\mu = 5$ for light sources placed at a radial distance of 10 units. This demonstrates that our formulation enables photometric stereo even for shadowed regions and non-parallel-ray light sources.

3) Comparison to other methods: Let us compare our shape recovery approach to other methods which have not necessarily been designed to handle shadows. We therefore allow all irradiance values and retain the data in double precision floating point format without first converting to 8-bit grayscale. The point light sources are positioned 3 units from the focal point, which, as we have mentioned, is in contrast to the classic assumption of distant light sources. The reasoning behind this placement is to demonstrate the deformation of the shape recovered by other methods when the model of nearby light sources, is not taken into account. We compare our approach with [1] and [8]. The surface normals provided to these two methods are first computed by assuming distant parallel ray light sources where the light direction is computed by averaging between the light source position and an approximate distance from the object (as done in [18]). [1] requires the use of full boundary conditions whereas [8] assumes so called natural boundary conditions which requires no additional information. Fig. 5 illustrates that the reconstruction by the other two methods becomes considerably worse when we add attenuation of $\mu = 1$, whereas our result remains consistent with an RMSE of 1.95E-2 units.

B. Real cases

In the last scenario we use a low cost endoscopic camera synchronized to four individually controlled low power white SMD light sources.

The images were taken in a dark environment and a calibration image was acquired without any lights activated. This was then subtracted from all subsequent images to account for ambient lighting. Any pixels with a value less than 20 were marked as shadow. This is a very simple heuristic but useful for our setup of multiple light sources. The intrinsic camera parameters were found using the Bouguet calibration toolbox [2] and all images had the effect of lens distortion removed. The depth to a single point on the object was measured manually. A more practical version could use a small laser pointer calibrated to the camera to extract an initial seed depth but we have not implemented this in these experiments. The processed images for each object can be seen in Fig. 6 with a rendering of a novel view of the same object based on our reconstruction. It is interesting to note that our method was able to successfully reconstruct the thumb and the ridges of the thumb-prints are clearly visible in the 3D render. Furthermore, despite the specularity and non-uniform albedo present in the steps images, the reconstruction is unaffected and preserves the straight lines of the edges of the steps. The hand model is also interesting because it illustrates that our method can handle real world scenarios in which there is a level of noise in the captured images.

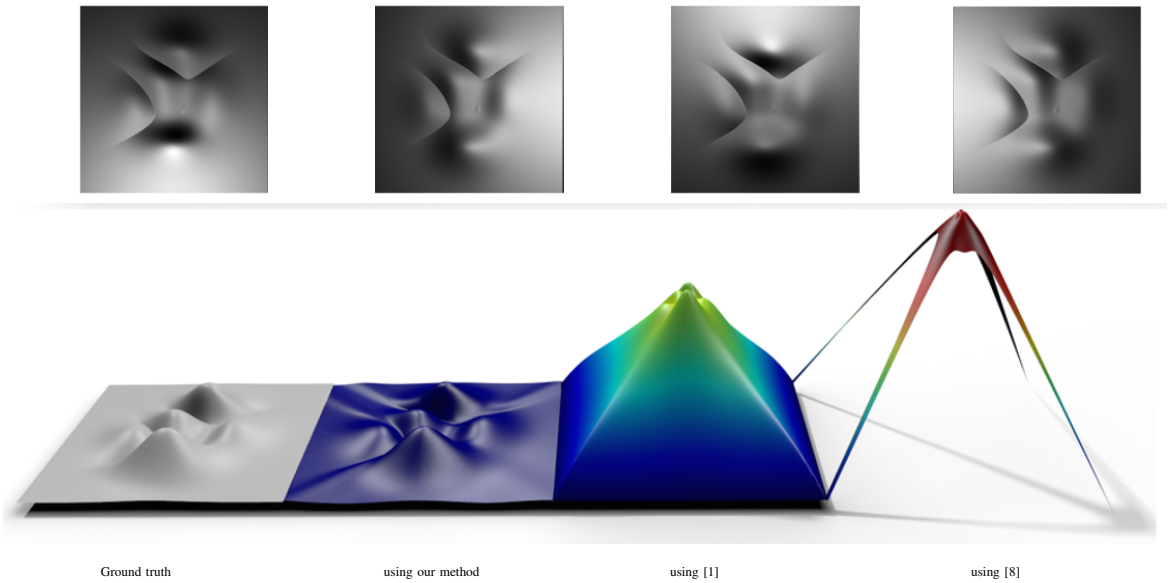

Figure 5. **First row** Images from AbsPeaks with $\mu = 1$ and lights positioned 3 units from the focal point. **Second row** From left to right: ground truth shape, reconstruction using our method $MSE = 3.82\text{E-}4 units^2$, the reconstruction using [1] $MSE = 5.19 units^2$, the reconstruction using [8] $MSE = 21.5 units^2$. All reconstructed shapes have their error maps textured on.

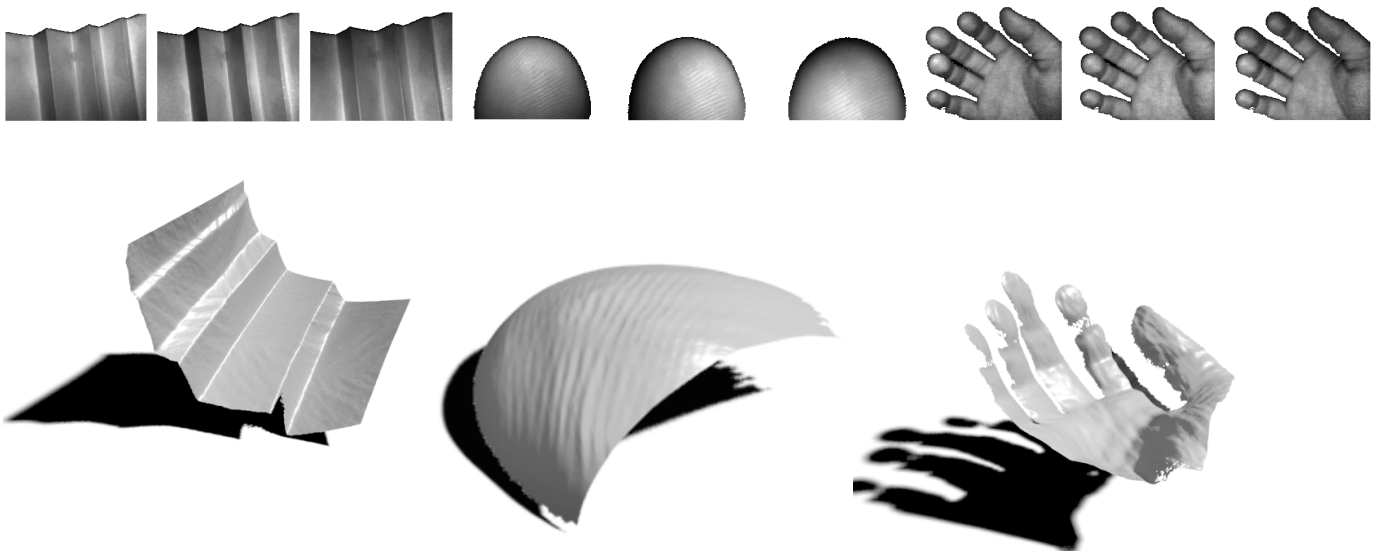

Figure 6. Novel views from reconstructions of real objects imaged with the experimental setup described in Section (VI-B). The captured images are shown on top of each reconstruction. The distance between the object and the scanner is: Model stairs 120mm, Thumb-tip 35mm, Hand 220mm

VII. Conclusions

An efficient model for shape reconstruction from the EPPS$_n$ problem was proposed. The main goal of the proposed model is to define a new differential formulation based on a quasi-linear PDE, where the well-posedness holds even in the presence of images with missing parts. We have shown that PDEs provide a strong way of modeling near-field photometric stereo that can be used to approximate completely general lighting scenarios. The model we presented overcomes the limitations of more classical approaches since the lighting is modeled realistically as a nearby source with fully general illumination. As far as we can tell this currently represents state of the art in physical modeling for any PS method. We are currently investigating even more general models for non-linearities such as specular effects and general BRDF functions. Our current experiments demonstrate that the model gracefully deals with non-linear light attenuations as well as non-uniform surface albedo and missing data. Furthermore, the suggested method handles real world surfaces and produces

quantitatively faithful surface reconstructions for nearby objects. The method is highly parallelizeable and future work will attempt an implementation on a kilo-core Graphics Processing Unit to demonstrate that real time endoscopic shape from PS is possible. Current ongoing work includes shrinking the scanning head and performing tracking so that reconstructed depth maps can be fused into a larger surface. Another issue is that of obtaining at least one known ground truth point (assuming a single connected surface). This has not been implemented but we are investigating using a laser dot or line calibrated to the camera.

ACKNOWLEDGMENT

This research was supported by European Community's FP7- ERC program, grant agreement no. 267414 and by Broadcom foundation. The authors would also like to thank Ron Slossberg for sharing his technical expertise for the experimental setup.

REFERENCES

[1] A. Agrawal, R. Raskar, and R. Chellappa. What is the range of surface reconstructions from a gradient field? ECCV'06, pages 578–591, 2006.

[2] J. Bouguet. Matlab camera calibration toolbox. In *http://www.vision.caltech.edu/bouguetj*, 2008.

[3] M. Breuß, E. Cristiani, J. D. Durou, M. Falcone, and O. Vogel. Perspective shape from shading: Ambiguity analysis and numerical approximations. *SIAM J. Imaging Sciences*, 5(1):311–342, 2012.

[4] A. M. Bruckstein. On shape from shading. *Computer Vision, Graphic, and Image Processing*, 44(2):139–154, 1988.

[5] M. Chandraker, J. Bai, and R. Ramamoorthi. On differential photometric reconstruction for unknown, isotropic brdfs. *IEEE Transactions on PAMI*, 2012.

[6] M. K. Chandraker, S. Agarwal, and D. J. Kriegman. ShadowCuts: Photometric Stereo with Shadows. In *CVPR*, 2007.

[7] T. Collins and A. Bartoli. 3d reconstruction in laparoscopy with close-range photometric stereo. In *MICCAI 2012*, volume 7511, pages 634–642. 2012.

[8] J.-D. Durou, J.-F. Aujol, and F. Courteille. Integrating the normal field of a surface in the presence of discontinuities. In *EMMCVPR*, pages 261–273, 2009.

[9] R. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on PAMI*, 10(4):439–451, 1988.

[10] C. Hernández, G. Vogiatzis, and R. Cipolla. Shadows in three-source photometric stereo. In *ECCV*, pages 290–303, 2008.

[11] B. K. P. Horn. Obtaining shape from shading information. *The Psychology of Computer Vision, Winston, P. H. (Ed.)*, pages 115–155, 1975.

[12] B. K. P. Horn. Height and gradient from shading. *The International Journal of Computer Vision, Winston*, 5:37–75, 1990.

[13] R. Kimmel and A. M. Bruckstein. Tracking level sets by level sets: A method for solving the shape from shading problem. *CVIU*, 62(1):47–58, 1995.

[14] R. Mecca and M. Falcone. Uniqueness and approximation of a photometric shape-from-shading model. *SIAM Journal on Imaging Sciences*, 6(1):616–659, 2013.

[15] R. Mecca, A. Tankus, A. Wetzler, and A. Bruckstein. A direct differential approach to photometric stereo with perspective viewing. *SIAM Journal on Imaging Sciences*, 7(2):579–612, 2014.

[16] T. Okatani and K. Deguchi. Shape reconstruction from an endoscope image by shape from shading technique for a point light source at the projection center. *Computer Vision and Image Understanding*, 66(2):119–131, 1997.

[17] T. Papadhimitri and P. Favaro. A New Perspective on Uncalibrated Photometric Stereo. In *IEEE International Conference on CVPR*, 2013.

[18] V. Parot, D. Lim, G. Gonzalez, G. Traverso, N. S. Nishioka, B. J. Vakoc, and N. J. Durr. Photometric stereo endoscopy. *Journal of Biomedical Optics*, 18(7):076017–076017, 2013.

[19] E. Prados and O. D. Faugeras. Shape from shading: A well-posed problem? In *CVPR*, volume 2, pages 870–877, 2005.

[20] L. B. Wolff and E. Angelopoulou. 3-d stereo using photometric ratios. In *In Proc. European Conference on Computer Vision*, pages 247–258. Springer, 1994.

[21] C. Wu, S. G. Narasimhan, and B. Jaramaz. A multi-image shape-from-shading framework for near-lighting perspective endoscopes. *IJCV*, pages 211–228, February 2009.

[22] S. Yeung, H. Tsui, and A. Yim. Global shape from shading for an endoscope image. In C. Taylor and A. Colchester, editors, *MICCAI*, volume 1679, pages 318–327. 1999.